

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF WISCONSIN

RAYTHEON COMPANY,

Plaintiff,

v.

CRAY, INC.,

Defendant.

OPINION AND ORDER

18-cv-318-wmc

In this patent case, plaintiff Raytheon Company asserts infringement claims against defendant Cray, Inc., based on two software-related patents. This case was originally filed in the Eastern District of Texas. The Texas district court issued an order construing a number of patent claims. (Dkt. #263.) Now before this court is defendant's motion for summary judgment on three, basic grounds: (1) plaintiff's direct infringement claims fail as a matter of law or, in the alternative, plaintiff lacks sufficient evidence from which a reasonable jury could find direct infringement; (2) plaintiff's induced infringement claims fail as a matter of law or, in the alternative, no reasonable jury could find induced infringement; and (3) even if its claims proceed to trial, plaintiff should be barred from pursuing damages based on sales to the United States under 28 U.S.C. § 1498. (Dkt. #392.)

On January 10, 2019, the court held a hearing on defendant's motion for summary judgment. For the reasons that follow, the court will grant defendant's motion, finding that the undisputed record forecloses a finding of direct infringement as to both patents-in-suit. Specifically, the court concludes that: (1) the alleged infringing technology does not meet the implicit ordering requirement in the independent claims of the '714 patent,

which require “shutting down” after selection of the node and retrieval of the policy; and (2) the alleged infringing technology does not meet the “plurality of hosts each executable at any of the nodes” limitation in the ‘274 patent because of the presence of service nodes. These findings, in turn, moot defendant’s challenges to plaintiff’s induced infringement claims, as well as plaintiff’s damages theory.

In addition, plaintiff Raytheon seeks summary judgment on defendant Cray’s counterclaims for unjust enrichment and conversion.¹ (Dkt. #388.) For the reasons explained below, the court concludes that Cray has failed to put forth evidence to support a finding that Raytheon used any confidential information its employee Ballew obtained or accessed through his role on an external review committee. Accordingly, the court will also grant plaintiff’s motion for summary judgment, entering judgment in its favor on these two, state law counterclaims.

UNDISPUTED FACTS²

A. The Parties and Origins of This Lawsuit

Plaintiff Raytheon Company is a Delaware corporation with its principal place of business in Waltham, Massachusetts. Defendant Cray, Inc., is a Washington corporation with its principal place of business in Seattle, Washington.

¹ During the briefing, the parties filed a joint stipulation of dismissal with prejudice of Cray’s counterclaim for breach of contract (dkt. #415), which the court accepts.

² Unless specifically noted, the court finds the following facts materials and undisputed. Moreover, because the parties’ proposed findings of fact were so extensive, this section only provides an overview of key facts. Other, undisputed facts relevant to particular claims and counterclaims are addressed in context in the opinion section.

Raytheon's licensing representative, Charles Neuenschwander of International Patent Licensing Co., sent Cray a letter dated March 20, 2015, asserting that Cray was infringing Raytheon's U.S. Patent No. 7,475,274 ("the '274 patent") and U.S. Patent No. 8,190,714 ("the '714 patent"), among others. There is no evidence that Cray was aware of these patents before receiving that letter.

On July 15, 2015, Cray filed an action in the Western District of Washington seeking a declaratory judgment of non-infringement as to nine of the patents Raytheon accused Cray of infringing, including the '274 and '714 patents. On September 25, Raytheon granted Cray a covenant not to sue as to five of the nine patents. That same day, Raytheon filed a complaint in the Eastern District of Texas alleging that Cray infringed the other four patents, the '274 patent and the '714 patent, along with two patents that are no longer part of this lawsuit -- U.S. Patent No. 8,335,909 ("the '909 patent") and U.S. Patent No. 9,037,833 ("the '833 patent"). This case progressed with the filing of infringement contentions and amendments to the same, an order on claims construction, serving of expert reports and extensive summary judgment briefing.

At that point, however, Cray challenged whether venue was proper in the Eastern District of Texas based on the Supreme Court's narrowing of the law of venue for *patent* cases in *TC Heartland LLC v. Kraft Foods Group Brands LLC*, 137 S. Ct. 1514 (2017). On June 29, 2017, the Texas district court issued an order finding that venue to be proper. Cray then filed a petition for writ of mandamus seeking reversal of the district court's decision and a transfer of the lawsuit to this court. On September 21, 2017, the Federal Circuit granted Cray's petition, directing the Texas court to transfer the case to an

appropriate venue. On April 5, 2018, the Texas court ordered that the case be transferred to this court. The transfer became official on May 2, 2018.

After the transfer, Raytheon withdrew its claims relating to the '909 and '833 patents, and the court directed the parties to submit summary judgment briefs, limited to those issues already raised in their motions submitted in Texas, scheduled a hearing on summary judgment, and set a trial date.

B. The Patents-in-Suit and Claims Construction

1. The '714 Patent

The '714 patent is entitled “System and Method for Computer Cluster Virtualization Using Dynamic Boot Images and Virtual Disk.” (Compl., Ex. B (“the 714 patent”) (dkt. #2-1).) At least in part, the '714 patent is directed at systems and methods for on-demand “provisioning” of computers with a “cluster” to manage and run “distributed applications” on the cluster. The patent defines “provisioning” as “the process of instantiating compute resources to the enterprise application by copying the local disk from a repository to the resource.” ('714 patent at 1:29-32.) The patent describes that typically provisioning takes over ten minutes, but the claimed invention can reduce that time to 15 seconds or less. (*Id.* at 1:29-34, 50-54.) A “cluster” is a group of computers that may be connected by a network coordinated by software. A “distributed application” or “job” is a piece of software created to address a large computing task that is split into pieces so it can be solved in parallel on many nodes working together. The parties define “node” as a “computing device.”

Before Raytheon’s submission of the '714 patent application, Cray does not dispute,

at least materially, that typically High-Performance Computing (“HPC”) Systems ran on dedicated computing resources within a system. Often, different applications required different computer environments, and switching an application from one computing device to another could require time-consuming resetting of the environment. Cray points out, however, that the concept of “provisioning” was already known at the time of the ’714 patent application. (Def.’s Resp. to Pl.’s Add’l PFOFs (dkt. #450) ¶ 6.)

The ’714 patent is also directed at software that will select available nodes from the cluster to run a distributed application, automatically shut down a selected node, and load that node with the operating system data to make it compatible. In doing so, the patent purports to disclose an efficient system for managing computers in a cluster by reducing the time required to ready each “node” to run a portion of a distributed application or job.

Raytheon alleges that Cray supercomputers infringe independent claims 1 and 29, as well as dependent claims 2, 4-6, 9, 11-14, 30, 32-34, 37 and 39-42 of the ’714 patent.

Independent claim 1 discloses a method comprising:

- selecting a distributed application;
- retrieving a policy associated with the distributed application;
- dynamically selecting one of a plurality of nodes;
- resetting a boot image of the selected node based at least in part on the retrieved policy,
 - wherein the boot image being compatible with the distributed application,
 - wherein resetting the boot image of the selected node comprises:
 - automatically shutting down the selected node;
 - resetting the boot image of the selected node; and
 - restarting the selected node using the reset boot image;
- associating a virtual disk image with the selected node based at least in part on the retrieved policy; and
- executing at least a portion of the distributed application on

the selected node, as reset, using the virtual disk image associated with the selected node, the execution performed by at least one processor of the selected node.

(’714 patent at 9:34-10:4.) Independent claim 29 contains a number of overlapping terms, but as opposed to a method, discloses a system comprising:

a plurality of nodes, each node comprising at least one processor; and
a management node communicably coupled to the plurality of nodes, the management node operable to:
select a distributed application; retrieve a policy with the distributed application; dynamically select one of a plurality of nodes;
reset a boot image of the selected node based at least in part on the retrieved policy, wherein the boot image being compatible with the distributed application, wherein, to reset the boot image of the selected node, the management node is operable to: automatically shut down the selected node; reset the boot image of the selected node; and restart the selected node using the reset boot image;
associate a virtual disk image with the selected node based at least in part on the retrieved policy; and
execute at least a portion of the distributed application on the selected node, as reset, using the virtual disk image associated with the selected node.

(’714 patent at 12:25-43.)

The Texas district court previously construed the following terms in the ’714 patent:

- “nodes” means “computing devices”
- “distributed application” means “a single instance of an application running across more than one node”
- “boot image” means “operating system data that is used to initialize a node”

- “[resetting / reset] a boot image of the selected node” has its plain meaning³
- “policy associated with the distributed application” means “one or more parameters that define the required characteristics of an execution environment in order to run the distributed application”
- “dynamically [selecting / select] one of a plurality of nodes” to mean “[selecting/select] a node at least in part at run-time based on one or more variables”
- “[associating / associate] a virtual disk image with the selected node” to mean “[providing/provide] the selected node with a pointer or other reference to a virtual disk drive”

(6/12/17 Order (dkt. #263) 9, 28, 38, 41, 45, 48, 49.)

2. The '274 Patent

The '274 patent is entitled “Fault Tolerance and Recovery in a High-Performance Computing (HPC) System.” (Compl., Ex. A (“’274 patent”) (dkt. #1-1).) The '274 patent is directed to handling node failures -- referred to as faults -- in an HPC system. The patent describes a system with a plurality of nodes, monitoring the status of the nodes for the occurrence of faults in executing hosts to run jobs and, in the event of a fault, substituting a free node in the system for the faulty node by discontinuing operation of the faulty node and booting the host at the free node.

Ordinarily, when a computer experiences a hardware failure, software and stored data remain unavailable until the failure has been resolved. Addressing this failure may require someone to identify manually where the error occurred and attempt to reroute the work to functioning nodes. Without fault tolerance, a large job involving many nodes

³ The Texas court made clear in the hearing on claims construction that this term “requires that the [node’s] boot image be reset based at least in part on the policy” (Aug. 4, 2016, Hr’g Tr. (dkt. #83) 90.)

might simply terminate and sit idle until someone intervened to restart the job on working nodes. Cray does not materially dispute any of this, although it asserts that even before the '274 patent, high-end supercomputers already included features providing for “fault intolerance operation.” (Def.’s Resp. to Pl.’s Add’l PFOFs (dkt. #450) ¶ 1.) Regardless, the '274 patent purports to improve upon prior art systems by dynamically recovering from job faults, which is achieved by a manager (a part of the software) monitoring the nodes to identify faults, and in the event of a fault, provisioning a free node with the proper execution environment and operating system (i.e., the same one running on the failed node) and transferring the pending job to the new node for completion.

Raytheon alleges that Cray supercomputers infringes independent claims 1, 16, 27 and 38, as well as dependent claims 2-8, 10, 12-15, 17-23, 25, 28-34 and 36. Independent claim 1 of the '274 patent provides:

A system for fault tolerance and recovery in a high-performance computing (HPC) system, the system for fault tolerance and recovery comprising;
a fabric coupling a plurality of nodes in an HPC system to each other, each node comprising a switching fabric integrated to a card and at least two processors integrated to the card;
storage coupled to the fabric and accessible to each of the nodes, the storage operable to store a plurality of hosts each executable at any of the nodes; and
a manager coupled to the fabric, the manager operable to monitor a currently running node in the HPC system executing a host and, if a fault occurs at the currently running node, discontinue operation of the currently running node and boot the host at a free node in the HPC system from the storage.

('274 patent at 69:16-31.)

Independent claim 16 recites:

A method for fault tolerance and recovery in a high-

performance computing (HPC) system, the method comprising:

monitoring a currently running node in an HPC system comprising a plurality of nodes, a fabric coupling the plurality of nodes of each other and coupling the plurality of nodes to a storage accessible to each of the plurality of nodes and operable to store a plurality of hosts each executable at any of the plurality of nodes, each node comprising a switching fabric integrated to a card and at least two processors integrated to the card; and

if a fault occurs at the currently running node:

discontinuing operation of the currently running node;

and

booting a host at a free node in the HC system from the storage.

(*Id.* at 70:19-34.)

Independent claim 27 provides:

One or more computer-readable storage media storing logic for fault tolerance and recovery in a high-performance computing (HPC) system, the logic when executed operable to:

monitor a currently running node in an HPC system comprising a plurality of nodes, a fabric coupling the plurality of nodes to each other and coupling the plurality of nodes to a storage accessible to each of the plurality of nodes and operable to store a plurality of hosts each executable at any of the plurality of nodes, each node comprising a switching fabric integrated to a card and at least two processors integrated to the card; and

if fault occurs at the currently running node:

discontinue operation of the currently running node; and

boot a host at a free node in the HPC system from the storage.

(*Id.* at 71:4-22.)

Finally, independent claim 38 recites:

A system for fault tolerance and recovery in high-performance computing (HPC) system, the system for fault tolerance and

recovery comprising computer-readable storage media comprising:

means for monitoring a currently running node in an HPC system comprising a plurality of nodes, a fabric coupling the plurality of nodes to each other and coupling the plurality of nodes to storage accessible to each of the plurality of nodes and operable to store a plurality of hosts each executable at any of the plurality of nodes, each node comprising a switching fabric integrated to a card and at least two processors integrated to the card; and

means for, if a fault occurs at the currently running node:

discontinuing operation of the currently running node; and

booting the host at a free node in the HPC system from the storage.

(*Id.* at 72:25-42.)

The parties agree that the claim elements Cray challenges in its motion for summary judgment are effectively the same for the independent claims, though claims 1, 27 and 38 are system claims and claim 16 is a method claim.

In its claims construction order, the Texas court construed the following terms:

- “integrated to” means “mechanically and electrically affixed to”
- “switching fabric” means “a network of switches and links between switches”
- “host” means “execution environment, including an operating system”
- “a manager coupled to a fabric, the manager operable to . . . from the storage” has its plain meaning

(6/12/17 Order (dkt. #263) 15, 32, 37, 59.)

C. The Accused Systems

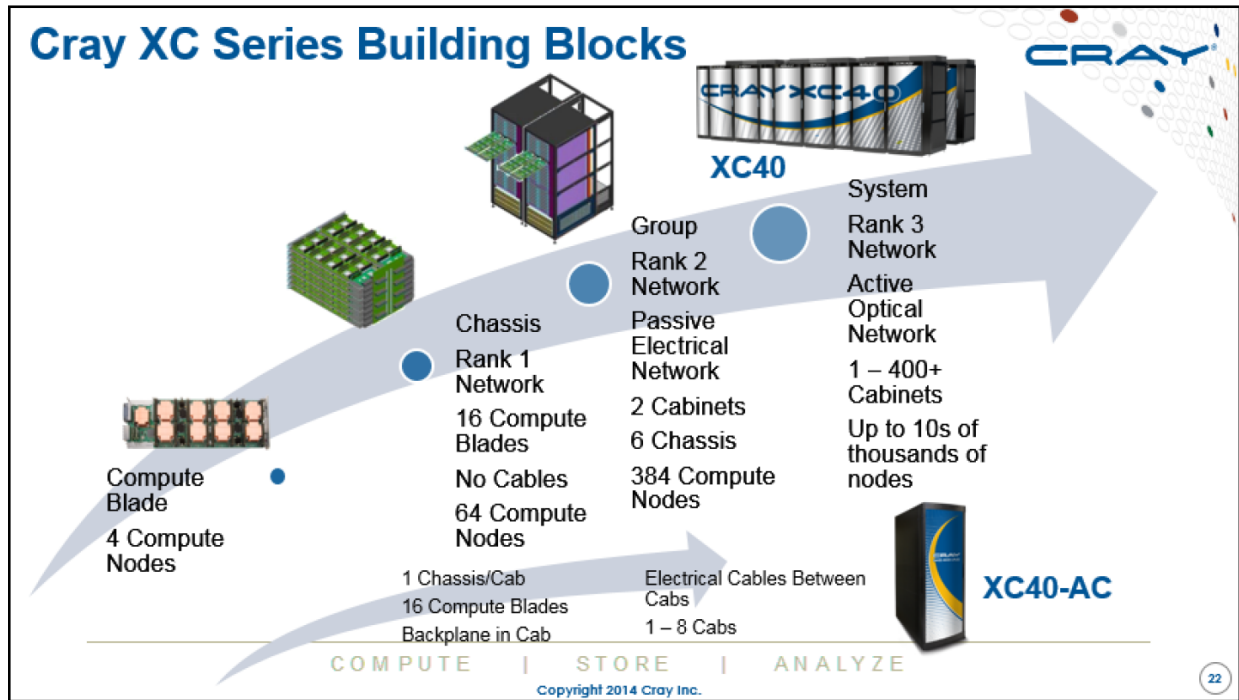
When combined with one of three kinds of third-party software -- PBS Professional,

Moab or Slurm -- Raytheon claims that three, specific lines of Cray supercomputers infringe the two patents-in-suit: the Cray XE, YK and XC systems (collectively the “Accused Systems”).

1. Hardware Architecture

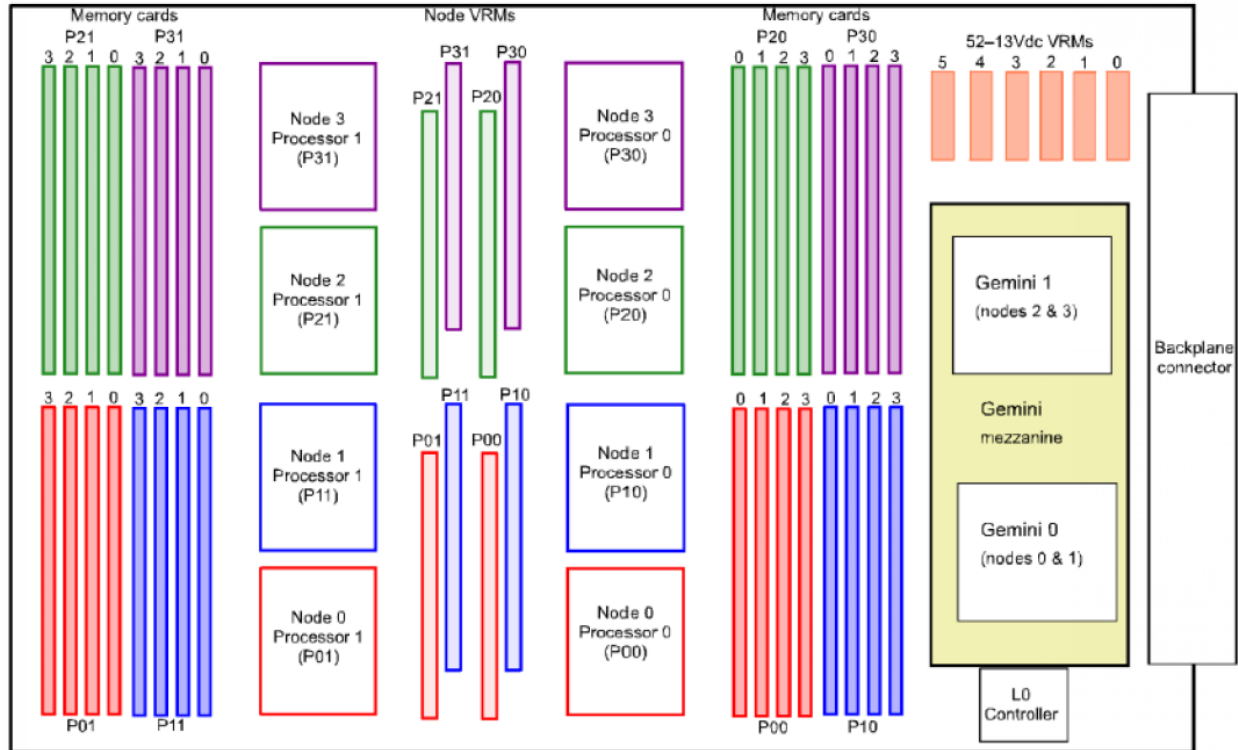
The Accused Systems are all large supercomputers, comprised of “nodes.” Material to the parties’ dispute, there are two types of nodes in the Accused systems: compute nodes and service nodes. Compute nodes comprise either one or two processors per node, whereas service nodes comprise only one processor per node. Regardless, each node acts as a processing unit within the system. Raytheon does not dispute these facts, but contends that the “plurality of nodes” term in the ’274 patent refers to “compute nodes” only, thus satisfying the patent limitation of “at least two processors integrated to the card.” (Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶ 98; Pl.’s Add’l PFOFs (dkt. #435) ¶¶ 19-25 (pointing to figures of grids and arguing from those figures that the “management nodes” excluded from the “plurality of nodes” are not “compute nodes”).) Cray, on the other hand, argues that the “plurality of nodes” encompasses all nodes.

The nodes communicate with other nodes through a custom Cray interconnect or “switch.” Collections of nodes and switches reside on “blades,” and collections of blades called “chassis” reside in cabinets. The image below shows the building blocks of a Cray XC system, including: (1) eight processors, comprising four nodes, placed on a metal frame known as a blade; (2) a set of 16 blades grouped together to form a chassis; (3) two cabinets, each containing three chassis; and (4) a row of multiple cabinets, comprising up to tens of thousands of processors.



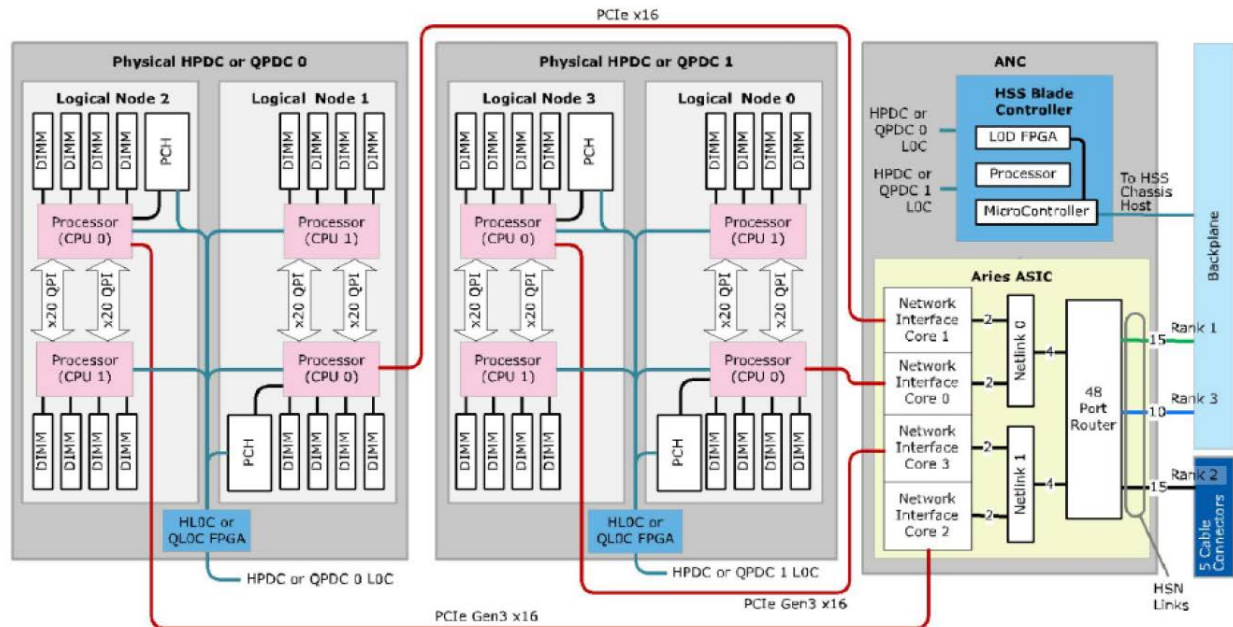
(Def.'s PFOFs (dkt. #394) ¶ 102.)

Cray designed the Accused Systems with the processors and switches affixed to different printed circuit boards or cards to facilitate simple repairs and easy upgrades. In the Cray XE and XK systems, the processors are affixed to a large printed circuit board, and the accused switch (known as a “Gemini” ASCI) is affixed to a separate printed circuit board (known as a “Gemini mezzanine card”), as depicted below in a figure showing a blade in the Cray XE System:



(*Id.* ¶ 108; *see also id.* ¶ 111 (displaying similar figure for Cray XK system).) Raytheon also does not dispute this, but points out that both of these boards are “mechanically and electrically affixed to a main board (or card),” thus satisfying another limitation of the ’214 patent. (Pl.’s Add’l PFOFs (dkt. #435) ¶¶ 11-12; *see also* Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶ 103.)

In the Cray XC systems, the processors are affixed to a smaller printed circuit boards (known as “processor daughter cards”), while the accused switch (known as “Aries” ASIC” is affixed to a separate printed circuit board called the Aires network card (or “ANC”), as depicted in the figure below:



(*Id.* ¶ 115.) Here, too, Raytheon does not dispute this set-up, but again points out that the ANC is “mechanically and electrically affixed to the same main board as two processors,” addressing Cray’s challenge to infringement of the ’214 patent. (Pl.’s Add’l PFOFS (dkt. #435) ¶ 14; *see also* Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶ 103.)

2. Software Architecture

Cray’s Accused Systems are tightly integrated, highly customized HPC systems that are custom-built for customers. Each of the Accused Systems uses a proprietary, customized version of the Linux operating system, and the Accused Systems only work with Cray’s custom Linux operating system.⁴ Service nodes run a version of Cray’s operating system called Cray Linux Environment (“CLE”), and compute nodes run a

⁴ Raytheon purports to dispute this last statement, but in doing so, only represents that a “CNL compute node can be configured with [Cluster Compatibility Mode].” (Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶ 130.) However, as discussed below, Raytheon admits that CCM is *not* an operating system. (*Id.* ¶ 142.)

version called Compute Node Linux (“CNL”). While a compute node may be allocated to act as a service node, and thus, run CLE, it is not possible for a service node to run CNL.

There are several other pieces of software potentially relevant to the parties’ dispute at summary judgment:

- System Management Workstation (“SMW”) is a terminal that a system administrator may use to boot and manage a Cray system. While the SMW remains active while the system is running, it plays no role in computation after the system is booted.
- Hardware Supervisory System (“HSS”) monitors system health metrics, and specifically monitors applications running at compute nodes, detecting if a node fails. When a node failure occurs, the HSS will receive a notification that the node’s “heartbeat” stopped. HSS connects to the nodes and to the SWM via an Ethernet network.
- Application Level Placement Scheduler (“ALPS”) serves as the interface between Cray’s internal system software and third-party software that may be installed on an Accused Systems, which means that it controls or limits what functions third-party software can perform on Cray’s machines.⁵
- Cluster Compatibility Mode (“CCM”) is a utility that translates messages that are written in a language other than Cray’s own message format into messages that a Cray operating system comprehends. As previously noted, there is no dispute that this is *not* an operating system, but the parties dispute whether it initializes or boots nodes. (Def.’s Resp. to PFOFs (dkt. #449) ¶ 142.)
- Cray Advanced Power Management and Control Utility (“CAPMC”) is a tool for monitoring power usage and performing power management operations within certain Accused Systems. While CAPMC may be used to limit the total amount of power used by a node, it does not perform that function automatically. Instead, CAPMC must be called on to do so by other software. The parties dispute whether third-party workload managers use CAPMC to power off or boot up a computer node. (Def.’s Reply to PFOFs (dkt. #449) ¶

⁵ The parties dispute the specific limitations ALPS places on third-party software, including whether third-party software is capable of powering down, booting or rebooting a Cray compute node or changing the boot image on a compute node. (Def.’s Reply to PFOFs (dkt. #449) ¶¶ 139-40; *see also* Pl.’s Add’l PFOFs (dkt. #435) ¶¶ 51-53, 56.)

146.)⁶

3. Workload Managers Software

Workload Managers Software (“WMS” or “workload managers”) provide tools to allow users to coordinate jobs on an HPC system. If there are many users vying to run jobs, a workload manager may queue incoming jobs until sufficient resources are available on the system to run them. As mentioned above, Raytheon’s infringement claims implicate three, third-party WMS: PBS Professional, Moab and Slurm (also referred to as SLURM). PBS Professional was created by Altair Engineering, Inc.; Moab was created by Adaptive Computing, Inc.; and Slurm was created by SchedMD, LLC.

Cray contends that these third-party workload managers are “restricted by design from administering the nodes within Cray’s systems.” (Def.’s PFOFs (dkt. #394) ¶ 152.) Raytheon disputes the extent of any restrictions, and whether they satisfy some of the claim limitations at issue in this case -- namely, powering off, rebooting a node and resetting a boot image. (Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶¶ 152-57.) In particular, Raytheon directs the court to its expert’s opinions and WMS manuals to support its argument that this software is capable of meeting those limitations. (Pl.’s Add’l PFOFs (dkt. #435) ¶¶ 31-47, 54 (describing features of PBS Professional, Moab and Slurm software).)

The Accused Systems are operational and able to run jobs without third-party WMS. Indeed, Cray ships some of its Accused Systems without WMS installed, with

⁶ In response to this and other proposed finding, Raytheon states that it disputes it and includes a cite to deposition testimony, often extensive citations across multiple pages, or to an exhibit, without providing any description of what is contained in the deposition transcript or exhibit. Including some description of the cited evidence, often in a parenthetical, is the typical practice and one the court appreciates.

Cray's customers selecting or installing WMS, should they wish, after purchasing an Accused System from Cray. Customers may work directly with a vendor to obtain the desired software, if any, or they may choose to purchase a license to WMS from Cray. In some cases, Cray also resells the WMS to the customer. In addition to the three WMS listed above, Cray currently supports at least one other and possibly three other workload managers. Indeed, some Cray customers use as many as six workload managers, including the three implicated in this case. Cray does not instruct nor encourage its customers to use any particular workload manager over another. The particular workload manager selected is typically a matter of customer preference and/or price.

D. Facts Relevant to Cray's Unjust Enrichment and Conversion Counterclaims

In addition to Raytheon's patent infringement claims, defendant Cray also asserts unjust enrichment and conversion counterclaims based on Raytheon's employee James Ballew's involvement in 2002 and 2003 in the development of a supercomputer known as Red Storm for Sandia National Laboratories. As mentioned, Raytheon also moves for summary judgment on these two counterclaims.

1. Ballew's Involvement in Red Storm Project

In 2001, Sandia National Laboratories issued a public "Statement of Work," containing requirements for a supercomputer system. Sandia issued an amended, public Statement of Work in 2002. In September 2002, Cray entered a contract with Sandia Corporation on behalf of Sandia National Laboratories to create a supercomputer system called "Red Storm" for Sandia National Laboratories.

That same year, Sandia put together an external review committee to review the design of the Red Storm System. Among others, a Raytheon employee, James Ballew, served on that external review committee, commencing with his participation at an August 14, 2002, committee meeting. While each page of Cray's presentation at that meeting was marked "Cray Inc. Confidential," there is no dispute that the information on many of those pages were publicly available and not proprietary. For example, there is no dispute that the requirements for the Red Storm RFQ were not confidential, having been circulated to the public for bids. Still, Cray contends that "[s]ome information disclosed during the August 14, 2002 Red Storm meeting was *not* public information," as confirmed by: (1) contemporaneous notes from the meeting that state "[s]ome items in the talk are proprietary and should be protected as such," albeit without detailing which "items" fit this category; and (2) deposition testimony of the Sandia engineer, who led the Red Storm project, that companies' *responses* to the Statement of Work were deemed proprietary. (Def.'s Add'l PFOFS (dkt. #427) ¶ 17.)⁷

Ballew further attended Red Storm design meetings on December 17, 2002, in

⁷ As part of an invalidity challenge to the hardware patents in this lawsuit, both parties pursued contrary positions regarding the confidential nature of these meetings. Specifically, Cray represented that "as of the first external review committee meeting on August 14, 2002, Red Storm constitutes prior art under § 102(a) because it was disclosed to members of the external review committee who were *not* subject to confidentiality obligations." (Def.'s Resp. to Pl.'s PFOFs (dkt. #428) ¶ 8 (quoting Def.'s MSJ (dkt. #140) 5) (emphasis added).) Cray has also represented that Red Storm was offered for sale at least as early as September 2002. Also contrary to its current position, Raytheon previously argued that summary judgment as to invalidity was inappropriate because Cray had "fail[ed] to provide clearly and convincingly which, if any, parts of the Red Storm design were public," and in a sur-reply brief, further argued, "Cray's Red Storm documents include original confidentiality labels on them and/or were produced under the attorneys' eyes only confidentiality designation in this litigation." (Def.'s Add'l PFOFs (dkt. #427) ¶¶ 47-48 (citing Pl.'s Opp'n (dkt. #157) 5; Pl.'s Sur-Reply (dkt. #199) 1-2).)

Mendota Heights, Minnesota on February 27, 2003, in Seattle, Washington, and on September 18, 2003, also in Seattle.⁸ Ballew attended at least one quarterly review meeting of the Red Storm Design in January 2003 as well. (Def.'s Add'l PFOFs (dkt. #427) ¶¶ 30-37.) The parties dispute whether information presented at any of these meetings was confidential, although Cray describes various categories of confidential information involved in the Red Storm project, including: Cray's proof of concept, hardware manufacturing, and interconnection of the nodes in the Red Storm system, as well as the configuration of those components on one or more motherboards; communication between the nodes; system software; efforts to test its design; timelines and procedures; and plans to commercialize the Red Storm system as a line of Cray supercomputing systems, including market analysis. (Def.'s Add'l PFOFs (dkt. #427) ¶¶ 53-62.) However, Cray's evidence as to the content of this information comes largely from quarterly review presentations for which Cray has failed to demonstrate or offer any evidence that Ballew even attended, much less reviewed, although in fairness, Ballew acknowledged at his deposition that it was possible that he had seen the January 2003

⁸ Cray submits evidence that other meetings were also held in Chippewa Falls, Wisconsin, on September 3, 2002, Baltimore, Maryland, on November 18, 2002, and Albuquerque, New Mexico, on December 3, 2003. However, the record evidence does not support a finding that Ballew attended any of those meetings. Moreover, as explained below, the relevance of the December 3, 2003, meeting is unclear, since it postdates the first non-disclosure agreement that Raytheon signed, and therefore is covered by the now-dismissed claim for breach of contract. In addition, Cray contends that Ballew attended the April, July and October 2003 quarterly review, but fails to point to any evidence in support. To the contrary, the deposition testimony to which Cray points in support is merely Ballew acknowledging that he remembered attending his first design meeting in Albuquerque in August 2002, as well as a couple more meetings in Seattle, consistent with the February 2003 and September 2003 dates. Even as to the January 2003 quarterly review, Ballew testified that it was *possible* he had seen the presentation before, but did not know. (Ballew Dep. (dkt. #384) 221-222.)

presentation before.

Moreover, it is undisputed that Sandia created a website called “QuickPlace” at which information relating to the Red Storm project could be posted. As a member of the external review committee, Ballew had been granted access to this website. Ballew also received emails notifying him of status reports and other materials posted to that website, although documents were not attached to the emails, and Cray fails to point to evidence that Ballew actually reviewed them online. At his deposition, Ballew further testified that he “really didn’t look at the QuickPoint [*sic*] site much” and that “[t]here was a lot of stuff I never downloaded. In fact, most of it.” (Pl.’s Resp. to Def.’s PFOFs (dkt. #447) ¶ 45 (citing Ballew Dep. (dkt. #384) 189, 202).)

Finally, Raytheon entered into multiple non-disclosures agreements (“NDAs”) with Cray, beginning on November 18, 2003.⁹

2. Evidence of Use of Red Storm Information by Raytheon

Before April 15, 2004, Raytheon had never developed or sold a supercomputer of its own design, although Raytheon points out that it has been involved in designing many aspects of supercomputers since the 1970s. In 2002, when Ballew was first invited to external review committee meetings for Cray’s Red Storm system, Raytheon acted as a “systems integrator,” which involves working between supercomputer manufacturers and customers to facilitate the design, purchase, delivery and installation of supercomputers.

⁹ Cray contends that during the August 14, 2002, meeting, inventory was taken to determine who had NDAs in place, and it was determined that Ballew and another individual did not. Nonetheless, there is no evidence that Ballew, individually or as a representative of Raytheon, signed an NDA at that time.


In that role, Raytheon had worked with Cray to sell supercomputers that were designed and built by Cray.

In 2003 and 2004, however, Raytheon attempted to develop and sell its own supercomputer called “Toro.” Ballew was Raytheon’s “technical leader” involved in the design of this supercomputer. In what appears to be an internal document dated March 4, 2004, Raytheon included Red Storm on a list of “Examples of Balance in HPC systems,” noting its node speed rate and link bandwidth. (Def.’s Add’l PFOFs (dkt. #427) ¶¶ 73-74.) There is no dispute that this information was not on the Red Storm public RFQ or Statement of Work, though Cray falls short of proving that this information was only obtainable through Ballew’s access to Cray’s confidential information. (Pl.’s Resp. to Def.’s Add’l PFOFs (dkt. #447) ¶ 76.)

In a presentation to a third-party dated September 2007, Raytheon represented, among other things, that a purchaser of its Toro architecture and patents, including its hardware patents, could: (1) “[g]ain more HPC sales and margins through features only available now on expensive proprietary systems like IBM’s BlueGene and Cray’s Red Storm”; (2) “[p]otentially block competitors from launching related products”; (3) “[g]ain an edge on HP, IBM, Dell, Cray, and Sun”; and (4) gain “BlueGene- and Red Storm-like functionality on standard blade servers that leverage open source software.” (Def.’s Add’l PFOFs (dkt. #427) ¶¶ 79-82; *see also id.* ¶¶ 83-85 (quoting similar language from October 2007 presentation).)

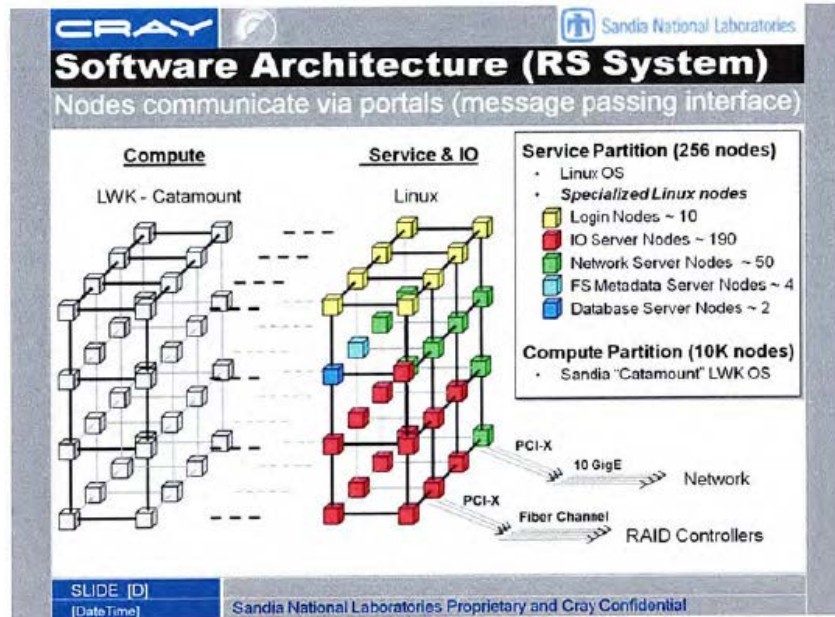
As described above, Raytheon obtained two hardware patents, the ’909 and ’833 patents. An application for these patents, U.S. Patent Application No. 10/824,874 (the

'874 application”), was filed on April 15, 2004. The patents have two named inventors: Ballew and Gary R. Early. Early did not begin to work on the design or the invention of these patent until January 2003. The '874 application contained the following figure:



(Def.’s Add’l PFOFs (dkt. #427) ¶ 88 (citing Mayer Decl., Ex. 52 (dkt. #436-52) 3.) The hardware patents contain that same image.

The Red Storm Project First Quarterly presentation, dated January 2003, which Ballew may have attended, contained the following slide:



(Def.'s Add'l PFOFs (dkt. #427) ¶ 90 (citing Mayer Decl., Ex. 13 (dkt. #426-13) 49).)

When presented with the figure on this slide at his deposition, Ballew acknowledged that it was familiar, that it was “all over the place,” and that he had “seen a drawing like that before.” (Ballew Dep. (dkt. #384) 222.) When presented with Figure 2A from the hardware patents, Ballew also testified that he had drawn that figure, that it was a “standard 3D Torus configuration” invented by “Euclid or somebody . . . before my time.” (*Id.* at 86.)¹⁰

Indeed, as Raytheon points out, Sandia’s public Statement of Work suggested this same 3-Dimensional configuration or topology as an option for the Red Storm project. (Pl.’s Resp. to Def.’s Add'l PFOFs (dkt. #447) ¶ 92 (citing McElDowney Dec., Ex. 2 (dkt. #390-2) 7).) Moreover, at his deposition, Robert Alverson, a Cray engineer working on the Red Storm project, could not identify any specific Cray employees’ work on Red Storm

¹⁰ See also “Torus Interconnect,” Wikipedia, https://en.wikipedia.org/wiki/Torus_interconnect.

that “fit into” or “fed into” the ’909 and ’833 patents. (Def.’s Resp. to Pl.’s PFOFs (dkt. #428) ¶ 35.)

OPINION

I. Cray’s Motion for Summary Judgment

In its motion for summary judgment, Cray seeks summary judgment on: (1) Raytheon’s claims of direct infringement of the two patents-in-suit; (2) Raytheon’s claims of induced infringement of the two patents-in-suit; and (3) Raytheon’s claim for damages covering sales of Accused Systems to the United States Government. Because the court concludes that the undisputed record forecloses a finding in Raytheon’s favor on the direct infringement claims, the court need not address the latter two bases for judgment in Cray’s favor.

A. ’714 Patent

Cray argues that the Accused Systems do not meet two limitations of the ’714 patent:

1. “resetting/ reset a boot image”; and
2. “shutting down the selected node”

(Def.’s Opening Br. (dkt. #393) 23.) As quoted in full above, these limitations are contained in part of claim 1, disclosing

A method comprising:

...

resetting a boot image of the selected node based at least in part on the retrieved policy, wherein the boot image being compatible with the distributed application, wherein *resetting the boot image* of the selected node comprises: automatically

shutting down the selected node; resetting the boot image of the selected node; and restarting the selected node using the reset boot image . . .

(’714 patent at 9:54, 59-65 (emphasis added).) Virtually the same language is used in independent system claim 29. As mentioned above, the Texas district court construed “boot image” to mean “operating system data that is used to initialize a node,” and assigned “[resetting / reset] a boot image of the selected node” its plain meaning.¹¹ While some of the arguments Cray posits are specific to one of the limitations, other arguments concern the relationship between these two limitations, including their required order in relation to other requirements. The court addresses those that are determinative below.

1. “Resetting/reset the boot image”

Raytheon contends that the Accused Systems have several features operable to reset the boot image. One way is through the Cluster Compatibility Mode software or CCM, which Raytheon contends the workload managers can invoke through a job script: “When a script invokes CCM, the Accused System will initialize a node (or nodes) with operating system data for the node to run the next job application.” (Pl.’s Opp’n (dkt. #431) 27.) Raytheon also contends that the workload managers themselves satisfy the claim element, specifically pointing to PBS Pro’s “file staging” feature, which is a “way to specify which input files should be copied on the execution host before the job starts, and which output files should be copied off the execution host when it finishes.” (*Id.* (quoting Pl.’s Add’l PFOFs (dkt. #435) ¶ 107).) Raytheon claims that Moab has a similar feature.

¹¹ The parties did not ask, and, thus, the court did not construe the term “shutting down.”

As Cray points out, however, a job script would also be required to satisfy the “resetting a boot image” limitation by means of a third-party workload manager. (Def.’s PFOFs (dkt. #394) ¶¶ 183, 185-89.) In fairness, the parties’ dispute begins with the definition of a “job script” itself. Cray contends that it is “new software that provides new functionality that did not exist before being written by a user or a system administrator.” (Def.’s PFOFs (dkt. #394) ¶ 182.) Raytheon defines it as simply “a list of commands stored in a text file,” though its expert Dr. Scherson acknowledged at his deposition that these set of commands do not exist until the script is written. (Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶ 182 (citing Scherson Dep. (dkt. #366) 87); Def.’s Reply to Def.’s PFOFs (dkt. #449) ¶ 182 (citing Scherson Am. Dep. (dkt. #367) 243); *see also* Pl.’s Add’l PFOFs (dkt. #435) ¶ 81; Def.’s Resp. to Pl.’s Add’l PFOFs (dkt. #450) ¶ 81.)

As part of this definitional dispute, Cray further contends that Raytheon lacks evidence that these job scripts have actually been written. In support, Cray again points to the deposition testimony of Raytheon’s expert, Dr. Scherson, who repeatedly and unequivocally stated that he had not seen job scripts necessary to meet this limitation. (Def.’s PFOFs (dkt. #394) ¶ 193.) In response, Raytheon points to Scherson’s further testimony that the features are nonetheless “operable.” (Def.’s Resp. to Pl.’s PFOFs (dkt. #430) ¶ 193.) This response, however, simply concerns whether a job script accesses/activates features *or* modifies the Accused Systems. Critically, it does not address Cray’s position that there is *no* evidence that these features have been used.

Of course, as Raytheon points out, for *system* claims -- like independent claim 29 of the ’714 patent -- infringement does not require a showing that the infringing feature is

“activated or utilized in any way.” *Fantasy Sports Properties, Inc. v. Sportsline.com, Inc.*, 287 F.3d 1008, 1018 (Fed. Cir. 2002); *see also Silicon Graphics, Inc. v. ATI Technologies, Inc.*, 607 F.3d 784, 794 (Fed. Cir. 2010) (“an apparatus claim . . . is nonetheless infringed so long as the product is ‘designed in such a way as to enable a user of that [product] to utilize the function’”). As was the case in *Fantasy Sports Properties*, this holding makes sense in light of the language of independent system claim 29, which covers “a management node communicably coupled to the plurality of nodes, the management node *operable* to” perform the required steps. (’714 patent at 12:27-28 (emphasis added).)

However, the same cannot be said for Raytheon’s method claims, including independent claim 1 of the ’714 patent, because method claims are “not infringed unless all the steps are carried out.” *Limelight Networks, Inc. v. Akamai Techs., Inc.*, 572 U.S. 915, 921 (2014). (*See* Pl.’s Opp’n (dkt. #431) 30 n.6 (acknowledging that “Cray is correct that for method claims, direct infringement requires actually invoking the infringing functionality”).) Therefore, the court agrees with Cray that the lack of evidence of these job scripts having been written dooms claim 1, but this argument does not resolve claim 29, the independent systems claim.

As for independent systems claim 29, Cray argues that the Accused Systems do not satisfy the “resetting a boot image” limitation, which would require a “modification.” As to this argument, the parties’ dispute largely turns on whether a job script constitutes a modification or simply operates as a means to access or activate features of the Accused Systems. For an answer, both parties direct this court to Federal Circuit opinions concluding either that (1) the infringing feature constituted a modification, and therefore

the product was not infringing, *or* (2) the infringing feature simply needed to be activated and, therefore, the accused product was capable of infringing.

In support of its position, Cray directs the court to *Nazomi Communications, Inc. v. Nokia Corp.*, 739 F.3d 1339, 1345 (Fed. Cir. 2014), which considered whether hardware present on the accused devices, but not functional without certain software, was capable of infringing. The plaintiff argued that the installation of the software did not constitute a modification, but the Federal Circuit rejected that argument, finding “the purchase and installation of the JTEK software clearly constitutes a ‘modification’ of the accused products,” and affirmed the district court’s grant of summary judgment of noninfringement. *Id.* at 1345; *see also Telemac Cellular Corp. v. Topp Telecom, Inc.*, 247 F.3d 1316 (Fed. Cir. 2001) (affirming entry of summary judgment in the defendant’s favor, concluding that “due to a restriction built into the software program,” the infringing feature -- directly placing international calls and the calculation of charges associated with such calls -- was not included in the accused device); *High Tech Med. Instrumentation v. New Image Indus., Inc.*, 49 F.3d 1551, 1555-56 (Fed. Cir. 1995) (“a device does not infringe simply because it is possible to alter it in a way that would satisfy all of the limitations of a patent claim,” but recognizing that if “a device is designed to be altered,” then infringement is not precluded).

In opposition, Raytheon directs the court to Federal Circuit opinions upholding a finding of infringement, or at least agreeing a fact issue required the case to proceed to trial, because the alleged infringing feature was present in the device and could be activated or accessed. *See Fantasy Sports Properties*, 287 F.3d at 1118 (reversing the district court’s

finding of noninfringement, concluding that “although a user must activate the functions programmed into a piece of software by selecting these options, the user is only activating means that are *already present in the underlying software*”); *Silicon Graphics*, 607 F.3d at 794 (reversing a finding of noninfringement on summary judgment, explaining that the alleged infringing processor had the structural means for performing a claimed function, even if an operating system was required for the products to function); *Versata Software, Inc. v. SAP America, Inc.*, 717 F.3d 1255, 1261-62 (Fed. Cir. 2013) (upholding a jury’s finding that the accused software infringed because it contained source code “configured to implement these particular functions” without requiring any modification).

This case appears to fall somewhere in the middle of the Federal Circuit’s previous decisions, largely depending on Raytheon’s shifting theory of infringement. If Raytheon primarily relies on the functionality of third-party workload managers, as it did in its infringement expert’s initial report, then Cray’s argument that resetting the boot image constitutes a modification to the Accused Systems has more traction. On the other hand, if Raytheon is primarily relying on Cray’s functionality to fulfill this limitation -- accessed by means of job scripts submitted through the third-party managers -- then its infringement theory appears to fall on the other side of this line of cases, simply requiring activation or accessing of a built-in feature. (*See Hr’g Tr. (dkt. #462)* 48-49 (counsel for Cray conceding that arguments raised with respect to job scripts are not implicated if Cray is relying on “aprun or a single command inside of [the Accused System]”).) As such, there appears a genuine factual dispute as to the appropriate characterization of the job script in each of Raytheon’s theories and, relatedly, whether a “modification” is required. Thus, summary

judgment is unavailable on this basis.

2. “Shutting down”

In its opposition brief, Raytheon sets forth its theory of how the “shutting down” limitation is met in the Accused Systems:

[U]sers submit jobs through job batch scripts using PBS Pro, Moab, or SLURM on Cray’s systems. When a job or application finishes running on Cray’s system, the system automatically cleans up the processes and shuts down the nodes. As described in Cray’s manuals, “[w]hen an application exits, either normally or due to error, ALPS must ensure that all resources allocated for the application are surrendered.” “Cleanup” then begins with “ALPS contacting each assigned compute node and sending a SIGKILL signal to any remaining application processes.” The system also automatically invokes a feature called the “Node Health Checker (NHC),” which in turns calls the “Compute Node Cleanup (CNCU).” The CNCU “return[s] compute nodes to the pool of available nodes with as much free memory as they have when they are first booted.”

(Pl.’s Opp’n (dkt. #431) 26 (internal citations omitted).)¹²

Cray offers several grounds to reject this theory of infringement. *First*, Cray argues that the theory rests on a “Cray-only” approach previously waived by Raytheon, rather than being dependent on the role of third-party workload managers or WMS. (6/1/18 Order (dkt. #347) 6-7 (finding that Cray waived any amendment to infringement contentions based on a “Cray-only” theory).) At first glance, this argument appears to lack

¹² In its opening brief on summary judgment, Cray addressed an alternative infringement theory premised on CAPMC fulfilling this step, but in response to Cray’s proposed findings of facts, Raytheon states that “CAPMC is not necessary for infringement,” apparently abandoning any theory premised on CAPMC shutting down and restarting nodes. (See Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶ 185.)

merit since the above description identifies a role for third-party WMS: submitting “job batch scripts using PBS Pro, Moab, or SLURM on Cray’s systems.” (Pl.’s Opp’n (dkt. #431) 26.) As became apparent during the summary judgment hearing, however, that shorthand reference to the waived theory of infringement as “Cray-only” may have been a bit of a misnomer. (*See* Hr’g Tr. (dkt. #462) 66.) Instead, for purposes of determining waiver, the proper inquiry is whether the theory Raytheon presses in its opposition motion was disclosed timely in Raytheon’s original report, and was only expanded upon in an amended report, or was it new and different from the originally disclosed theories of infringement. If the latter, then Raytheon waived those new theories for reasons explained in this court’s June 1, 2018, Order.

In particular, the theory on which Raytheon now relies in opposing summary judgment argues that features in the Accused Systems themselves meet the patent limitations like ALPS, even though third-party workload managers are mentioned. To be fair, ALPS is discussed in the original report of plaintiff’s expert, Dr. Isaac D. Scherson (Scherson Orig. Rept., Ex. A (dkt. #134-2) ¶¶ 23-28), but ALPS is simply an interface or conduit for other software programs, whether internal to Cray’s system or a third-party software, including the third-party managers. More critically for purposes of summary judgment, Raytheon does *not* reply on ALPS as the feature that would fulfill the shutting down function; instead, Raytheon’s counsel argued during the hearing, it is the Node Health Checker that performs the shutting down function under this new theory of infringement. (Hr’g Tr. (dkt. #462) 24 (“The Node Health Checker is . . . controlled by ALPS, and it cleans up the node and returns it to an initial state.”); *id.* at 27 (again

explaining that Node Health Checker fulfills the shut-down step).)

This is where waiver becomes relevant, since Dr. Scherson's original report does *not* describe or otherwise disclose *any* theory of shutting down a node based on the Node Health Checker feature. Indeed, the original expert report does not mention this feature *at all*, except in a single instance in the middle of a multi-page reproduction of materials from a Cray Brochure, and even then, only in a discussion of the '274 patent. (Scherson Orig. Rept. (dkt. #134-2) p.326 ("Node Health Checker – When Node Health Checker (NHC) runs in Normal Mode, the xtcheckhealth binary is automatically invoked by ALPS upon the termination of an application.").) Rather, Scherson describes in his original report how the provisioning features in the third-party workload managers satisfies the shutting down limitation in element claim 29. (*Id.* at pp.193-98.) Thus, Raytheon's basis for opposing Cray's motion for summary judgment as to this limitation rests on a waived theory of infringement.

Second, even putting aside the waiver argument, the undisputed facts support Cray's argument on the merits that, at least with respect to the Accused Systems at issue in this case, "a job/application can be assigned only to already-booted (up and running) nodes, and those nodes remain active and powered on until completion of the job." (Def.'s Opening Br. (dkt. #393) 24-25; *see also* Def.'s PFOFs (dkt. #394) ¶ 132, 152-57.) Moreover, the workload managers are restricted in their actions. (*Id.*) Specifically, they cannot power off or restart a node. (*Id.*) In response, Raytheon counters that "shutting down" does not mean "powering off;" rather, shutting down simply means "the action of putting a node in an initial state" or "initializing." (Pl.'s Opp'n (dkt. #431) 29.) Once

again, Cray argues that Raytheon's theory rests on a "belated and waived claim construction." (Def.'s Reply (dkt. #448) 25 (citing *Bettcher Indus., Inc. v. Bunzl USA, Inc.*, 661 F.3d 629, 640-41 (Fed. Cir. 2011) (finding no abuse of discretion when district court refused to introduce new claims construction after "ample opportunity" to do so)).) Given the parties' apparent dispute over the meaning of the term "shutting down," both sides could be faulted for failing to seek claims construction for this term, as well as for constructions of the broader three-step process -- shutting down, resetting the boot image, and starting the node.

Ultimately, however, Cray acknowledged during the hearing that a complete powering down is not required under the patent; instead, the term could be construed fairly as a "warm boot." (Hr'g Tr. (dkt. #462) 64 ("A warm boot is not something we take issue with.")). Cray also maintains that the Accused Systems do not perform a "warm boot" either, but the court need not reach this issue for the reason that follows.

3. Required Order

Even if the two limitations -- both shutting down and resetting the boot image -- are met, Cray contends that the limitations do not occur in the required order. Specifically, Cray argues that the patent discloses "a three-part process/function: (1) automatically shutting down the node; (2) resetting the boot image of the node; and then (3) restarting the node," all of which must occur *after* selection of the node and *before* the execution of a job. (Def.'s Opening Br. (dkt. #393) 24.) Tellingly, Raytheon's response to this argument is relegated to a footnote:

To the extent Cray is arguing the claims require a particular

sequence of steps (Br. at 19), that is not required by the patent. ‘714 patent, 9:37-44 (“[M]any of the steps in this flowchart may take simultaneously and/or in different orders than as shown. Moreover, system 100 may use methods with additional steps, fewer steps, and/or different steps, so long as the methods remain appropriate.”) It is also contrary to the law. *Altiris, Inc. v. Symantec Corp.*, 318 F.3d 1363, 1369-70 (Fed. Cir. 2003) (“Unless the steps of a method actually recite an order the steps are not ordinarily construed to require one.”).

(Pl.’s Opp’n (dkt. #431) 26-27 n.5.)

There are two core flaws with this response. First, the quoted language from the specification, which contemplates the possibility of “different orders,” cannot broaden or otherwise augment the language of the claims themselves. See *ACTV Inc. v. Walt Disney Co.*, 346 F.3d 1082, 1088 (Fed. Cir. 2003) (“[T]he analytical focus of claim construction must begin, and remain centered, on the language of the claims themselves.”). Second, the caselaw does not foreclose an *implicit* required order for system claims and, therefore, Raytheon’s contention that the Cray’s position is “contrary to the law” is overly simplistic at best.

The Federal Circuit instructs that the “first step” in determining whether an order is required is “to look at the claim language to determine if, as a matter of logic or grammar, they must be performed in the order written.” *Altiris*, 318 F.3d at 1369-70. The Federal Circuit further explains that words like “the” and “said” generally “refer[] back to the same term recited earlier in the claim.” See *Wi-LAN, Inc. v. Apple, Inc.*, 811 F.3d 455, 462 (Fed. Cir. 2016). As such, the court in *Wi-LAN* concluded that “[b]ecause [the term] ‘the modulated data symbols’ refers back to these already-randomized symbols, the claims impose the disputed ordering requirement.” *Id.* at 462-63.

Relying on independent claim 29, Cray convincingly explains why the following order, with some expressly noted flexibility, is required:

- “select a distributed application” must come first, because following terms refer to “*the* distributed application;”
- “retrieve a policy associated with *the* distributed application” must come after select a distributed application but before the execution step since retrieval of a policy is a necessary step to executing the distributed application;
- “dynamically select one of a plurality of nodes” must occur before any term referring to “*the* selected node,” including the steps of resetting a boot image, associating a virtual disk image and execution;
- “reset a boot image of *the* selected node based at least in part on *the* retrieved policy” must occur after selection of the node and retrieval of the policy but before execution as described below;
- “wherein *the* boot image being compatible with *the* distributed application” provides a condition of “reset a boot image” that occurs in that same step;
- “wherein, to reset *the* boot image of the selected node, the management node is operable to: automatically shut down *the* selected node; reset *the* boot image of the selected node; and restart *the* selected node using *the* reset boot image” also provides a condition of “reset a boot image” with sub-steps, which necessarily require that any restarting occurs after shutting down the node;
- “associate a virtual disk image with *the* selected node based at least in part on *the* retrieved policy” must occur after the selection of the node and the retrieved policy; and
- “execute at least a portion of *the* distributed application on *the* selected node, *as reset*, using *the* virtual disk image associated with the selected node” must occur after selection of the distributed application, selection of the selected node, resetting of the boot image on that selected node, and association of a virtual disk image; in other words, this step must be last.

(See '714 patent at 12:29-43; Cray's Reply (dkt. #448) 17-20.)

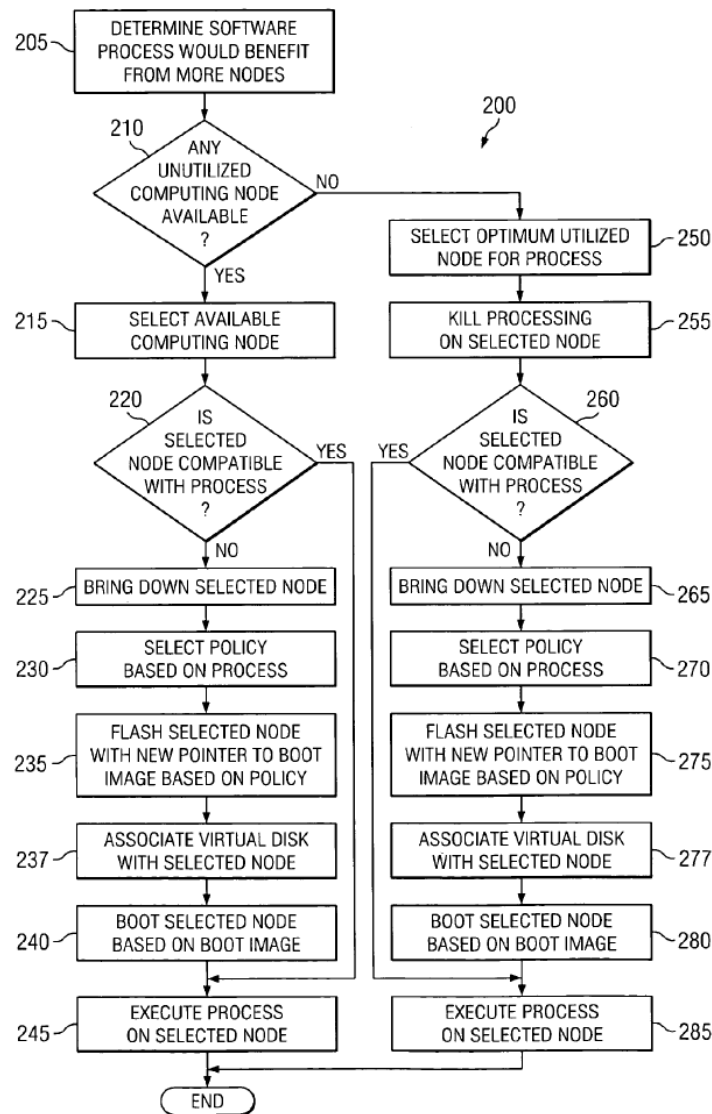
As for the Accused Systems, Cray then argues that Raytheon has put forth *no* theory of infringement that satisfies the required order of shutting down and resetting the boot image *after* selection of the node and *before* execution. Instead, under Raytheon's theory,

the shutting down of a node occurs *after* execution and *before* selection of the node. As an initial matter, Cray's argument as to the relationship between the shutting down and execution steps is a bit of a red herring. In considering jobs or processing cycles, Cray defines the start of a cycle as execution, but this line drawing is artificial; the start of the cycle could just as easily be defined as the step immediately following execution.

The critical issue in terms of required ordering is whether resetting the boot image, including shutting down, happens *before* the selection of the node. However, Raytheon's theory of infringement describes the Accused Systems generally cleaning nodes after execution, readying them for another execution, without *any* reference to a node being selected with a specific distributed application in mind, and then before shutting down the selected node, resetting that node with a boot image, and finally restarting that node, all as required by the claim language. (Hr'g Tr. (dkt. #462) 123-24 (“[I]n our view[,] the shutdown happens at the end of a job.”).)

At the hearing, perhaps recognizing that their response to the ordering argument was underwhelming, Raytheon's counsel also pointed to Figure 2 in the '719 patent, which is a “flowchart illustrating an example method **200** for dynamically rebooting a node **108** within one embodiment of the present disclosure”:

FIG. 2



('714 patent at Fig. 2; *id.* at 8:18-20.)

As counsel pointed out, this flowchart contemplates bringing down the selected node *before* selection of the policy based on the application, identified above at steps 225 and 230, respectively. However, this embodiment *still* cabins the order of the invention: a node is selected, brought down, and rebooted with reference to a distributed application, even if the specific policy is selected after the node is brought down. In other words, the

embodiment does *not* support Raytheon's argument that the patent requires no ordering and is satisfied by simply bringing a node back to its initial state following execution of a job. While Figure 2 suggests some greater flexibility than that described in Cray's ordering of the steps, both independent claims of the '714 patent require a definite order in the context of executing a specific application process: selection of a node, shutting down of that node, resetting of the boot image, and restarting of that node. Here, the court agrees with Cray that to hold otherwise both defies logic and is contrary to the claimed invention itself.

In contrast, there is no dispute on summary judgment that Raytheon's infringement theory simply rests on features of the Accused Technology, which initialize nodes after execution of an application, readying them for another job, without any reference to a specific application or job. Indeed, Raytheon conceded as much during the hearing. (*See* Hr'g Tr. (dkt. #462) 123-24.) Accordingly, the court concludes that Raytheon has failed to raise a genuine issue of material fact that the Accused Technology satisfies the required order in the independent claims of the '714 patent. The court, therefore, will grant summary judgment to Cray on Raytheon's claims of infringement as to that patent.

B. '274 Patent

As for the '274 patent, defendant argues that plaintiff cannot demonstrate that the Accused Systems meet the following three limitations:

- “boot the host as a free node”
- “storage operable to a plurality of hosts each executable at any of the nodes”
and

- “each node comprising a switching fabric integrated to a card and at least two processors integrated to the card.”

(Def.’s Opening Br. (dkt. #393) 33.) As quoted above, each of these limitations is present in all four of the independent claims of the ’274 patent.

1. “Boot the host at a free node”

Raytheon argues that this limitation is met in the Accused Systems in at least one of two ways. First, and principally, Raytheon contends that a user can trigger the functionality by running a batch job on the Accused systems using Slurm, PBS Pro or Moab. “For PBA Pro and Moab, ‘the basic process for creating and running batch jobs is to create a job script that includes aprun commands, then use the qsub command to run the script.’” (Pl.’s Opp’n (dkt. #431) 17 (quoting Pl.’s Add’l PFOFs (dkt. #435) ¶¶ 27-28).) While Slurm uses an “sbatch,” rather than a qsub, command to submit job scripts, it, too, includes aprun commands. (*Id.* (citing Pl.’s Add’l PFOFs (dkt. #435) ¶ 29).) Raytheon goes on to explain the “aprun” command functions meet the “boot the host at a free node” limitation because they include both: “automatically attempt[ing] to relaunch a job on healthy, available nodes (i.e., excluding the failed node)”; and “[s]et[ting] an environment variable on the compute nodes.” (*Id.* at 18.) Second, and with little explanation, Raytheon contends each of the three workload managers have “requeue” functions that also satisfy this limitation. (*Id.* at 19.)

With respect to the first theory at least, Cray responds with the same modification argument described above as to the ’714 patent: plaintiff’s theory of infringement requires writing a job script, which in turn is a modification of the Accused Systems, and, therefore,

cannot constitute infringement. For the same reasons described above, this argument is difficult to untangle factually and legally. Thankfully, the court again need not reach its merits in light of stronger, clearer-cut bases for summary judgment.

Cray also argues that the “aprun” theory of infringement falls within the waived “Cray-only” theory. Here, too, Raytheon’s theory appears to implicate the workload managers by using them to submit job scripts. Still, as Cray points out, any substantive discussion of the “aprun” commands only appears in Dr. Scherson’s supplemental expert report, addressing proposed amended infringement contentions that the court previously found had been waived. While aprun is also mentioned generally in Scherson’s original report in the overview of system operation section, and specifically in his discussion of Cray’s node monitoring and node failure management (*see* Scherson Original Rept. (dkt. #134-2) ¶ 31), Scherson’s specific discussion of how the alleged technology meets the specific limitation of “boot the host at a free node” only focuses on certain features of the workload managers, without any reference to “aprun” commands. (*Id.* at ¶¶ 110-28.)¹³ Arguably, this very limited, timely disclosure of the role of the workload managers in submitting jobs scripts is not enough to pull it within the properly disclosed infringement contentions. Instead, in his amended expert report, Scherson sets forth an actual infringement theory based on the aprun command in the accused instrumentalities.

Finally, why Cray’s waiver argument does not apply to Raytheon’s theory of

¹³ Indeed, in this section of his report, Dr. Scherson provides a data dump of excerpts from various documents from the three workload managers, without drilling down as to how aspects of those documents meet the specific limitations of the asserted claim.

infringement based on the “requeue” by third-party workload managers software or WMS, as Raytheon pointed out at the hearing, its explanation of the basis for that theory is exceedingly thin. Indeed, Raytheon simply points to manuals and deposition testimony of individuals familiar with the workload managers as evidence that all three include a feature allowing the workload manager to “requeue a job in the event that it detects a node failure.” (See, e.g., Pl.’s Add’l PFOFs (dkt. #435) ¶ 68; see also Hr’g Tr. (dkt. #462) 11 (repeating same requeue argument at hearing); Pl.’s Presentation at slide 10 (“PBS Pro will detect if a node fails when a job is running on it, and will automatically requeue and schedule the job to run elsewhere.”).) Still, Raytheon stops short of explaining how this function fulfills the limitation of “booting the host on a free node.”¹⁴ Because the alleged infringing technology does not meet the second challenged limitation, addressed directly below, the court will not definitively decide this issue either.

2. “Plurality of hosts each executable at any of the nodes”

Cray next argues that it is entitled to a finding of noninfringement of the asserted claims of the ’274 patent because the limitation requiring a “plurality of hosts each executable at any of the nodes” is not met. As described above, the Accused Systems all contain compute nodes and service nodes. Compute nodes are capable of running both

¹⁴ Cray argues that its nodes are similarly not “booted nor rebooted when a job/application is selected, and thus are not rebooted based on the retrieved policy associated with the selected job,” nor do the Accused Systems “‘automatically shut down’ nodes after job selection and prior to job execution.” (Def.’s Opening Br. (dkt. #393) 24.) In response, Raytheon argues that the claim does not require that the *node* is booted; rather, the claim simply requires booting the *host* at the node, not the node itself. Raytheon’s argument appears to have merit, but Cray contends in its reply that this is a “distinction without a difference” because booting the host at the node necessarily requires booting the node. (Def.’s Reply (dkt. #448) 31.) The court also need not resolve this dispute given plaintiff’s failure to show the Accused Systems satisfy the second limitation.

CNL and CLE, which the parties agree are the “hosts,” previously defined by the Texas district court as an “execution environment, including an operating system.” In contrast, service nodes are designed to run *only* CLE. (Pl.’s Resp. to Def.’s PFOFs (dkt. #430) ¶ 222 (not disputing that service nodes cannot execute CNL).) From this, Cray argues that the plurality of hosts -- here CNL and CLE -- are not “*each* executable at any of the nodes.” Specifically, CNL is not executable at the service nodes. (Def.’s Opening Br. (dkt. #393) 37-38.)¹⁵

As a factual matter, Raytheon concedes that service nodes in the Accused Systems cannot execute CNL, but instead argues that this independent claim limitation applies only to compute nodes, essentially as a matter of claim construction. (Pl.’s Opp’n (dkt. #431) 15.) In support, Raytheon points to: (1) language from the specification that supposedly contemplates a “plurality of nodes” being limited to compute nodes; and (2) other language in the patent that suggest “management” nodes fall outside of that limitation.¹⁶ Specifically, Raytheon points to the description of a “a grid 110 comprising a *plurality of nodes* 115.” (Pl.’s Opp’n (dkt. #431) 15 (citing ’274 patent at 4:19) (emphasis added).) From this, Raytheon concludes the plurality of nodes, labeled as 115, must be the “‘computing nodes’ that the patented invention is concerned with monitoring for faults as

¹⁵ Moreover, “service nodes” would also not fulfill the “at least two processors” limitation since it is undisputed that service nodes only have one processor, but for reasons that are not clear, the parties do not focus on this limitation in their briefing.

¹⁶ Raytheon’s argument is confusing because the claim describes both a “plurality of *nodes*” and a “plurality of *hosts*.” Consistent with the parties’ treatment of this language, the court concludes that the “each executable at *any* node” language is limited to the “plurality of nodes” disclosed in the independent claims.

they run job applications.” (*Id.* (citing ’274 patent at 3:11-14 (“System 100 provides users with HPC functionality dynamically allocated among various *computing nodes* 115.”) (emphasis added))).) In contrast, Raytheon concludes, the ’274 patent’s description of other “management nodes” are “physically or logically separated from the plurality of HPC nodes.” (*Id.* at 15-16 (citing ’274 patent at 4:39; *id.*, fig. 1 (depicting management node 105 as separate from the grid 110 comprising the plurality of nodes).)

Setting aside its renewed waiver argument,¹⁷ Cray challenges the merits of Raytheon’s proposed construction of the “plurality of nodes” limitation. As an initial matter, Cray points to the Federal Circuit’s decision in *Apple Inc. v. Samsung Electronics Company*, 695 F.3d 1370, 1379 (Fed. Cir. 2012), but this case does not resolve, nor even assist in resolving, the parties’ dispute. In *Apple*, the court considered an argument that a limitation requiring a “plurality of modules . . . wherein . . . each heuristic module corresponds to a respective area of search and employs a different, predetermined heuristic argument” was satisfied as long as there is “one such ‘one plurality,’” regardless of whether there were “other modules and heuristic algorithms” included. *Id.* at 1378-79. The Federal Circuit rejected this argument, explaining,

Apple’s argument essentially urges us to hold that “plurality” refers not to all but a subset of modules. As we pointed out, however, the district court has construed “plurality” to mean “at least two,” without any indication that the term refers to a hand-picked selection of a larger set. . . . Accordingly, despite

¹⁷ In its reply brief, Cray again argues that Raytheon waived any claims construction for this term by failing to seek it timely before the Texas district court. This court, however, is reluctant to find waiver and leave the meaning of this term undefined, whether at summary judgment or at trial. If anything, the parties’ dispute about timeliness reinforces the benefits of construing claims in conjunction with summary judgment motions.

the use of “comprising,” claim 6 is not amenable to the addition of other modules that do not use different heuristic algorithms because such addition would impermissibly wipe out the express limitation that requires every module to have a unique heuristic algorithm.

Id. at 1379. Here, however, Raytheon does *not adopt* Apple’s position that the “plurality of hosts each executable at any nodes” is met even though there are service nodes *within* the definition of “plurality of nodes”; rather, it is challenging the meaning of “plurality of nodes,” arguing that the term as used in the patent-in-suit in Apple did not include service nodes. As such, the parties’ dispute appears to turn solely on claims construction, and whether the patent’s use of “plurality of nodes” is limited to compute nodes in the context of the ‘274 patent.

Turning specifically to the language of the patent-in-suit here, Cray points to other language in the specification describing nodes 115 broadly as “any computer or processing device such as, for example, blades, general-purposes personal computers (PC), Macintoshes, workstations, Unix-based computers or any other suitable devices.” (Def.’s Reply (dkt. #448) 36 (citing ’274 patent at 3:62-66).) Cray also points to Figure 5A, with a grid of nodes labeled 115, which are *not* limited to “compute nodes.” (*Id.* (citing ’274 patent at p.8 (also identifying external i/o nodes, fs servers, fs metadata service and database servers as 115).)

In response to Raytheon’s position that a reference to “management node 105” is somehow akin to service nodes in the Accused Systems in particular, Cray further points to the opinion of Raytheon’s own expert, Dr. Scherson, in the context of the ’714 patent: “A management node need not be a single physical node, but may be distributed

throughout many nodes, including service nodes, compute nodes, and I/O nodes.” (*Id.* (citing Scherson Rept. (dkt. #377) A-26 ¶ 61)).) Cray also sets forth in great detail other reasons why its service nodes do not fall within the definition of a “management node,” including that service nodes are “part of the same ‘grid’ and the same interconnect[ed] network in Cray’s accused systems,” and therefore are not “physically and logically separated from the compute nodes.” (Def.’s Resp. to Pl.’s Add’l PFOFs (dkt. #24).) Indeed, during the hearing, counsel for Raytheon conceded that service nodes are different from management nodes. (Hr’g Tr. (dkt. #462) 108; *see also id.* at 70 (discussing Scherson’s testimony that Cray’s “SMW” is the management node identified as 105 in the patent).)

From all of this, the court agrees with Cray’s position that the “plurality of hosts each executable at any of the nodes” is not limited to the compute nodes. Having construed the patent term this way, the court further concludes that summary judgment of noninfringement of the ’274 patent claims is warranted since there is no dispute that only one host is executable at each service node in the Accused Systems.

3. “Integrated to” the same card

In light of the court’s finding as to the “plurality of hosts each executable at any of the nodes” limitation, the court need not decide the merits of Cray’s argument with respect to a third limitation requiring an integration to the same card. Instead, the court will comment briefly on this last challenge. Cray contends that the Accused Systems do not infringe the ’274 patent claims because the switching fabric and processors are not “integrated to” the same card. As set forth and displayed in the graphics above, each of

the Accused Systems contain the processors on one circuit board and the switch (or switching fabric) on another. The parties' dispute turns on whether this arrangement satisfies the "integrated to" the same card limitation as construed by the Texas district court.

The Texas court dealt extensively with the meaning of this term, concluding that "integrated to" means "mechanically and electrically affixed to." (6/12/17 Order (dkt. #263) 15.) In so holding, the court expressly rejected Raytheon's broader proposed construction that "integrated to" means "mechanically and electrically connect[ed] to." Moreover, the court did so out of concern that Raytheon's construction would allow for too many intervening layers of connections between the nodes, switching fabric, processors and card for the limitation to have meaning. The Texas court also rejected Cray's narrower construction, limiting the term to "fabricated on the same motherboard without the use of daughterboards or other components." (Def.'s PFOFs (dkt. #394) ¶ 93; Pl.'s Resp. to Def.'s PFOFs (dkt. #430) ¶ 93.)

Under the court's construction, therefore, the parties' dispute constitutes a classic factual dispute: whether the "daughter" or "mezzanine" cards on which the switch or switching fabric is mounted in the Accused Systems is sufficient to find that the processors and switching fabric are "mechanically and electrically affixed to" the same board. Nevertheless, having also concluded that plaintiff's infringement claim fails because the accused technology does not meet the "plurality of hosts each executable at any of the nodes" limitation, this is no longer a *material* factual dispute. Accordingly, the court will grant judgment in defendant's favor on plaintiff's infringement claims of the '274 patent.

II. Raytheon's Motion for Summary Judgment

Raytheon also seeks summary judgment on Cray's counterclaims for unjust enrichment and conversion. As described above, Cray alleges that Raytheon "misappropriated Cray's confidential and proprietary information, including the design of the Red Storm system, in obtaining the '909 and '833 patents," and that Raytheon used and benefited from the use of "Cray's confidential and proprietary information." (Def.'s Am. Ans. & Countercls. (dkt. #118) ¶¶ 120-21, 125).)

Before turning to Cray's proffered evidence as to the elements of these claims, the court needs to address three initial matters. First, as alluded to above and discussed at the hearing, Cray's now dismissed breach of contract claim defines the parameters of its unjust enrichment and conversion claims. The parties' non-disclosure agreements controls after the first was entered into on November 18, 2003. As such, Cray's tort claims are necessarily limited to information Ballew received about the Red Storm project before that date. (*See* Def.'s Opp'n (dkt. #425) 23 (acknowledging NDAs first signed on November 18, 2003, but stating that "the activity subject to Cray's unjust enrichment counterclaim began with Ballew's involvement in the Red Storm project starting more than one year earlier, in 2002").)¹⁸

Second, in its opposition brief, Cray implies that at least its unjust enrichment claim need not be limited to confidential and proprietary information (*See* Def.'s Opp'n (dkt.

¹⁸ Depending on its terms, a non-disclosure agreement may also dictate the remedies for unauthorized use of confidential information disclosed before its execution, but neither side has asserted that is the case here.

#425) 2 (arguing that Raytheon’s attempt to limit an unjust enrichment claim to confidential and proprietary information as a “legal standard” is in error).) More directly, Cray disputes Raytheon’s proposed findings “to the extent [it] suggests a legal requirement that Cray’s unjust enrichment claim requires confidentiality.” (*See, e.g.*, Def.’s Resp. to Pl.’s PFOFs (dkt. #328) ¶¶ 18, 19.) While the court agrees that an unjust enrichment claim need not be limited to confidential information as a matter of law, defendant’s alleged claim is so limited here. (*See* Def.’s Am. Answ. & Countercls. (dkt. #118) ¶¶ 120-21, 125 (pleading counterclaims based on Raytheon’s use and benefit from the use of “Cray’s confidential and proprietary information”).) Regardless, as explained below, this limit is immaterial in light of the court’s finding that Cray has failed to put forth sufficient evidence of use by Raytheon of Red Storm information, confidential or otherwise, in its development of the hardware patents or in any other way.

Third, the parties agree that Washington law governs Cray’s counterclaims. To prove unjust enrichment, therefore, Cray must prove “a benefit conferred upon [Raytheon] by [Cray]; an appreciation of knowledge by [Raytheon] of the benefit; and the acceptance or retention by [Raytheon] of the benefit under such circumstances as to make it inequitable for [Raytheon] to retain the benefit without the payment of its value.” *Young v. Young*, 191 P.3d 1258, 1262 (Wash. 2008) (en banc). As for the conversion claim, Cray must prove Raytheon “intentionally interfere[d] with chattel belonging to another, either by taking or unlawfully retaining it, thereby depriving the rightful owner of possession.” *Alhadeff v. Meridian on Bainbridge Island*, 220 P.3d 1214, 1223 (Wash. 2009).

As for Cray’s proffer, with regard to these elements, Raytheon’s motion for summary

judgment boils down to two assertions: (1) there is insufficient evidence from which a reasonable jury could conclude that Ballew obtained confidential information from Cray about the Red Storm project between August 2002 and November 2003 (the “relevant period”); and (2) there is insufficient evidence from which a reasonable jury could conclude that Raytheon used confidential information in the development of the Toro supercomputer or hardware patents or otherwise.

As for the first assertion, while Cray has advanced sufficient evidence to create an issue of fact as to whether *some* of the information provided to the external review committee was confidential and proprietary, a jury would be left to *speculate* as to what information Ballew actually accessed or reviewed during the relevant period. Specifically, as described in more detail above, Cray presents evidence that Ballew attended four meetings: the August 14, 2002, kick-off meeting; the December 17, 2002, meeting in Minnesota; a February 27, 2003, meeting in Seattle; and a September 17, 2003, also in Seattle. With respect to the August 14 meeting, while each page of the presentation was labeled confidential, there is no dispute that only “some items” in the presentation were confidential or proprietary, consistent with Cray’s position earlier in this case that Red Storm was a prior art reference in challenging the validity of Raytheon’s hardware patents.

When pressed to identify the *specific* confidential information, Cray only points to third-party responses to Sandia’s Statement of Work, but neither explains how these responses would constitute proprietary or confidential information belonging to *Cray*, thus serving as a basis for its conversion and unjust enrichment claims. Cray’s evidence of confidential information is similarly lacking with respect to the other three meetings. At

the December 17, 2002, meeting, Cray identifies disclosure of the “Seastar chip design” as confidential information to which Ballew was privy (Mayer Decl., Ex. 9 (dkt. #426-9)), but Cray does not offer proof as to what aspects of this design were confidential, nor does Cray present *any* evidence that the chip design was a benefit to Raytheon’s design of the Toro computing system or its hardware patents. For the other two meetings, plaintiff points even more amorphously to reports that summarized the “outcome of the above-mentioned review,” containing a number of bullet points, but again with no explanation as to what information was confidential, nor how Raytheon benefitted. (*Id.*, Ex. 10 (dkt. #426-10); *id.*, Ex. 11 (dkt. #426-11).)

Instead, Cray primarily relies on information presented during *quarterly reviews* in January, April, July and October 2003, but Cray fails to offer evidence that Ballew attended any of these meetings or later accessed the information presented. While Cray apparently believes that opportunity alone is enough for the jury to infer all the actual elements of a conversion or unjust enrichment claim, this limited proof instead invites rank speculation by the trier of fact. Even if Cray had advanced sufficient evidence to raise a genuine issue of material fact as to whether Ballew accessed confidential information, it failed to put forth sufficient evidence from which a reasonable jury could conclude that Raytheon actually *used* that information in the development of the Toro architecture and its patents. As detailed above, Cray’s evidence of use is limited to three documents: (1) a Raytheon document dated March 4, 2004, listing Red Storm on a list of “Examples of Balance in HPC Systems,” which notes node speed rate and link bandwidth; (2) a Raytheon presentation dated September 2007, discussing the competitive advantage of Raytheon’s

Toro architecture vis-à-vis Red Storm and other products; and (3) the Red Storm software architecture figure, which is similar to a figure in the hardware patents.

As for the source of the node speed and link bandwidth information, Cray points to a “Cray RS / Strider NDA Briefing” document dated November 2003 (Mayer Decl. Ex. 49 (dkt. #426-49)), but fails to demonstrate Ballew’s access. Indeed, there is *no* evidence that Ballew attended a meeting in or around November 2003, or even that there was a Red Storm external review meeting in November 2003. Moreover, given the date of the presentation, it appears likely that it was produced and shared with someone at Raytheon *after* the first NDA was signed on November 18, 2003, and thus would fall within the scope of the now dismissed breach of contract claim. In further support, the document itself is denominated as an “NDA Briefing.” Absent some evidence that this document came into Raytheon’s hands during the period of time relevant to Cray’s tort claims, there is no basis for a reasonable trier of fact to rely on it to find Raytheon obtained, much less benefitted from, this information.

Indeed, putting the unique issues with respect to this first document aside, all three documents fail to support a finding that Raytheon actually used confidential information, whether viewed individually or in combination. The first two documents constitute relatively generic competitive market information and posturing. While the node speed and link bandwidth for the Red Storm computer are listed on one Raytheon document, this information is provided alongside similar information for other competing computing systems. Moreover, Cray offers no evidence that Red Storm’s node speed and link

bandwidth information was held in confidence at that time.¹⁹ As for the similarity between the software architecture figure displayed in the November 2003 presentation and the Raytheon hardware patent application, the basic design is a common illustration of a 3D Torus (*see* '909 patent (dkt. #390-8) 6:66-7:3), which Ballew testified at his deposition traces back to Euclid (*see* Ballew Dep. (dkt. #384) 86). Indeed, even Sandia's original, public Statement of Work mentions this topology as an option for the Red Storm project. (McEldowney Dep., Ex. 2 (dkt. #390-2) 7.)

For all these reasons, the court concludes that Cray has failed to put forth sufficient evidence from which a reasonable jury could find that Raytheon used Cray's confidential information in developing the Toro computing system, in applying for the hardware patents, or in some other, unspecified manner. Accordingly, the court will grant summary judgment to Raytheon on Cray's unjust enrichment and conversion claims.²⁰

¹⁹ As opposed to the other products, Red Storm had not "launched" on the date of either documents creation. Accordingly, a trier of fact might infer that this information was not public at the time. However, Raytheon offers no proof that this was so. More importantly, Raytheon offers no proof those third-parties given access to this information were told either piece of information was being held in confidence for a product that would launch shortly.

²⁰ Raytheon also seeks summary judgment on the basis that both claims are time-barred under the applicable, three-year statute of limitations, arguing that Cray had, or at least should have, learned of these claims before November 2013, which is three years before the filing of the counterclaims. *See Seattle Prof's Eng'g Employees Ass'n v. Boeing Co.*, 991 P.2d 1126, 1133-34 (three-year statute of limitations for unjust enrichment claims), *opinion corrected on denial of reconsideration*, 1 P.3d 578 (Wash. 2000); *Hudson v. Condon*, 6 P.3d 615, 619 (Wash. App. 2000) (three-years statute of limitations for conversion claims). Specifically, Raytheon argues that Cray should have learned of its hardware patents, or the technology disclosed in them, either at a 2004 conference or as part of a 2012 acquisition of a third-party -- the latter having previously entered into a license agreement with Raytheon, which, in turn, listed one of Raytheon's hardware patents. (Def.'s Resp. to Pl.'s PFOFs (dkt. #428) ¶¶ 36-40.) However, Raytheon's statute of limitations defense turns on disputed fact issues, making it unavailable for purposes of summary judgment.

ORDER

IT IS ORDERED that:

- 1) Defendant Cray, Inc.'s motion for summary judgment (dkt. #392) is GRANTED. At the close of this case, the clerk's office is directed to enter judgment in favor of Cray on Raytheon's claims of infringement of U.S. Patent No. 7,475,274 and U.S. Patent No. 8,190,714.
- 2) Plaintiff Raytheon Company's motion for summary judgment (dkt. #388) is GRANTED. At the close of this case, the clerk's office is directed to enter judgment in favor of Raytheon on Cray's conversion and unjust enrichment counterclaims.
- 3) The parties' joint stipulation to dismiss with prejudice Cray's counterclaim for breach of contract (dkt. #415) is ACCEPTED. At the close of this case, the clerk's office is directed to enter judgment in favor of Raytheon on Cray's breach of contract counterclaim, dismissing that claim with prejudice.
- 4) On or before April 22, 2019, the parties are directed to file statements as to what, if any, issues remain.

Entered this 15th day of April, 2019.

BY THE COURT:

/s/

WILLIAM M. CONLEY
District Judge